

METHOD, SYSTEM, AND PROGRAM FOR ERROR
HANDLING IN A DUAL ADAPTOR SYSTEM
WHERE ONE ADAPTOR IS A MASTER

5

RELATED APPLICATIONS

[0001] This application is related to the copending and commonly assigned United States patent application entitled "Method, System, and Program for Error Handling in a Dual Adaptor System" having attorney docket no. TUC920010103US1, which patent application was filed on the same date herewith and is incorporated herein by reference in its entirety.

10

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0002] The present invention relates to a system, method, and program for error handling in a dual adaptor system.

15

2. Description of the Related Art

[0003] In a storage loop architecture, such as the Serial Storage Architecture (SSA), a plurality of disks are interconnected to one or more adaptors so that either of the adaptors can access the one or more loops of interconnected disks. An adaptor may include two or more ports to allow connection to one or more loops. For each loop on which the adaptor communicates, one adaptor port connects to a first disk in the loop and the other port connects to another disk in the loop. Additional adaptors may be added to the loop, such that one port on each other adaptor connects to one disk and another port connects to another disk so that the additional adaptors are placed within the loop. Additional details of the SSA architecture and different possible loop topologies are described in the International Business Machines Corporation (IBM) publication "Understanding SSA Subsystems in Your Environment", IBM document no. SG24-5750-00 (April, 2000), which publication is incorporated herein by reference in its entirety.

20

25

FOR OFFICIAL USE ONLY

[0004] One or more computer systems, such as storage subsystems, host system, etc., may include the adaptors connecting to the loop. Adaptors that share a loop must intercommunicate to coordinate accesses to disks in the shared loop. High end storage systems, such as the IBM Enterprise Storage Server (ESS), can detect errors in the ability of an adaptor in another system to communicate with the local operating system even though such detected adaptor is still capable of communicating on the network. In such instances, the system detecting the problem will delay I/O processing for a timeout period that corresponds to the time required for the other system including the adaptor to initiate an error recovery procedure. This timeout period must take into account all different timeout periods and error recovery procedures that could occur within the detected system unable to communicate with the adaptor. In many cases the timeout period can extend for several minutes.

[0005] In storage systems requiring high availability, such as storage systems for critical uses, any delays in I/O processing are generally unacceptable. Thus, extensive delays in I/O processing, such as a delay resulting from the lengthy timeout period for the error recovery process at the detected system, would be unacceptable in a high availability system.

[0006] In addition to delays that may result from having to wait for the system housing the other adaptor to reset, additional delays may be incurred when a master adaptor is subject to the reset. The master adaptor, which is the configurator with the highest unique identifier (ID), is responsible for configuring each port in the network with various parameters and coordinating the processing of asynchronous events such as dynamic changes in the network configuration. If a master adaptor is reset, then in the SSA architecture, the adaptor having the next highest unique identifier will be designated as the master. Following reassignment of the master node, each remaining adaptor on the loop adjusts internal routing algorithms under direction from the new master initiator, so that frames are automatically rerouted to avoid the break. This allows devices to be removed or added to the loop while the subsystem continues to operate without interruption.

[0007] Upon resetting an adaptor, the system will experience a brief I/O delay to coordinate the reset adaptor entering a disabled state. If a slave is reset, then the I/O delay may only be a few seconds. However, if the master is reset, then the I/O delay may double to 8 to 16 seconds due to the additional time needed to switch the master to
5 another adaptor.

[0008] For these reasons there is a need in the art to provide improved error handling that reduces timeout delays in systems where two adaptors are capable of accessing the storage devices and reduces delays associated with resetting the master adaptor.

10 SUMMARY OF THE PREFERRED EMBODIMENTS

[0009] Provided is a method, system, and program for processing Input/Output (I/O) requests to a storage network including at least one storage device and at least two adaptors, wherein each adaptor is capable of communicating I/O requests to the at least one storage device. An error is detected in a system including a first adaptor, wherein the
15 first adaptor is capable of communicating on the network after the error is detected. In response to detecting the error, a master switch timer is started that is less than a system timeout period if the first adaptor is the master. An error recovery procedure in the system including the first adaptor would be initiated after the system timeout period has expired. An operation is initiated to designate another adaptor in the storage network as
20 the master if the first adaptor is the master in response to detecting an expiration of the master switch timer.

[0010] In further implementations, a reset request is sent to the first adaptor after the master switch timer expires.

[0011] In still further implementations, in response to detecting the error, a monitoring
25 state is initiated to monitor I/O requests transmitted through a second adaptor. In response to receiving an I/O request, an I/O delay timer is started that is less than the system timeout period. A reset request is sent to the first adaptor in response to detecting an expiration of one started I/O delay timer.

[0012] Yet further, the steps of initiating a monitoring state, starting the I/O delay timer and sending the reset request may be performed by a device driver executing in an operating system.

[0013] The described implementations provide improved error handling techniques
5 when detecting a problem with the ability of another adaptor to communicate to the system in which the adaptor is housed. The described error handling techniques reduce delays in I/O processing when detecting errors in the ability of other adaptors to communicate to the system in which they are housed.

10

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

FIG. 1 illustrates a computing environment in which aspects of the invention are implemented; and

15

FIG. 2 illustrates an alternative computing environment in which aspects of the invention are also implemented; and

FIG. 3 illustrates logic implemented when detecting a problem in the system housing another adaptor in accordance with implementations of the invention; and

FIGs. 4a and 4b illustrate additional logic implemented when detecting a problem
20 in the system housing another adaptor in accordance with implementations of the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0015] In the following description, reference is made to the accompanying drawings
25 which form a part hereof and which illustrate several embodiments of the present invention. It is understood that other embodiments may be utilized and structural and operational changes may be made without departing from the scope of the present invention.

[0016] FIG. 1 illustrates a computing environment in which aspects of the invention are implemented. Computer systems 2a and 2b include at least one adaptor 4a and 4b to enable communication with a loop 6 of interconnected storage devices 8a, 8b...8j, such as hard disk drives, tape drives, optical disk drives, etc. Each adaptor 4a and 4b includes a plurality of ports 10a, 12a, 14a, 16a and 10b, 12b, 14b, and 16b, respectively. In FIG. 1, ports 10a, 10b, 12a, and 12b are shown as connected to the loop 6. The other ports 14a, 14b, 16a, and 16b may connect to additional loops. Further, FIG. 1 provides an example loop configuration. Those skilled in the art will appreciate that there are numerous loop configurations that may be implemented between one or more adaptors.

10 [0017] The systems 2a and 2b may comprise any computing device known in the art, such as a host system, storage subsystem, workstation, server, etc. Each system 2a and 2b includes an operating system 20a and 20b, respectively, and device drivers 22a and 22b that enable communication between the operating systems 20a, 20b and the adaptors 4a, 4b. Each of the storage devices 8a, 8b...8j would also include ports (not shown) to connect to the loop 6. In certain implementations, the loop 6 may be implemented using the Storage Serial Architecture (SSA) or any other protocol enabling multiple hosts to communicate with multiple storage device, such as Fibre Channel, Ethernet, Gigabyte Ethernet, Infiniband, etc. The loop 6 may be part of a Storage Area Network (SAN), Wide Area Network (WAN), Local Area Network (LAN), the Internet, an Intranet, etc.

15 [0018] FIG. 2 illustrates an alternative implementation where the adaptors 54a and 54b are housed in different Input/Output bays 52a and 52b within the same system 50. Each I/O bay 52a and 52b provides a separate power boundary. Thus, adaptors 54a and 54b are on separate power boundaries. Each adaptor 54a and 54b includes ports 60a, 62a, 64a, 66a and 60b, 62b, 64b, and 66b, respectively, for communication with one or more loops (not shown) according to any loop topology known in the art. In the implementation of FIG. 2, the system 50 includes one operating system 70 and separate device drivers 72a and 72b to manage the operations of adaptors 54a and 54b. Alternatively, only one instance of the device driver may execute to manage the operations of both the adaptors 54a and 54b. Each I/O bay 52a and 52b provides

different communication channels and power sources for the adaptors 54a and 54b, respectively.

[0019] FIG. 3 illustrates logic implemented in the device drivers 22a and 22b (FIG. 1), as well as device drivers 72a and 72b (FIG. 2) to handle a situation where an error is
5 detected in the system housing the other target adaptor 4b. For instance, the target adaptor 4b may be unable to communicate with the operating system 20b due to a problem in the bus interfacing between the adaptor 4b and the system 2b in which the operating system 20b is executing. For instance, the adaptor 4b may be implemented on a Peripheral Component Interconnect (PCI) card and the PCI bus between the system 2b
10 and adaptor 4b may be malfunctioning. In such cases, the target adaptor 4b itself is operational, but does not fully function because of another problem in the system 2b housing the target adaptor 4b.

[0020] The error handling process of FIG. 3 begins at block 100 where the device driver in one adaptor, e.g., device driver 22a in adaptor 4a, detects an error within the
15 system 2b including the target adaptor 4b, even though no error was detected at the target adaptor 4b. In such case, the source adaptor 4a would be able to continue processing I/Os against the storage devices 8a, 8b...8j because the target adaptor 4b is available on the loop 6. In response to detecting the error, the device driver 22a begins two concurrent (or non-concurrent) processes at blocks 102 and 106. At block 102, the
20 device driver 22a starts a window timeout timer 102. In certain implementations, the duration of the window timeout timer coincides with a time during which the system 2b including the target adaptor 4b is guaranteed to have completed an error recovery procedure. This window timeout timer includes the time that would pass before the other system 2b would initiate the error recovery procedure after detecting a problem. The
25 error recovery procedure initiated by the system 2b would likely eliminate the detected error. Upon the system 2b beginning the error recovery procedure by resetting or undergoing a power cycle, the target adaptor 4b would go offline and not participate in the loop 6. In certain implementations, each adaptor 4a, 4b on the loop is aware of all active adaptors. Once an adaptor becomes unavailable or unable to communicate on the

loop 6, then the adaptors remaining on the loop are notified of such state change and remove the unavailable adaptor from loop registries (not shown) each adaptor maintains that indicates available adaptors on the loop. The capability of one system to detect errors in other systems, such as problems a remote system may have in communicating with internal adaptors is known in the art of high availability storage subsystems, such as the IBM Enterprise Storage Servers (ESS).** Thus, if the device driver 22a, 22b takes no action for the window timeout period, the device driver 22a, 22b can be assured that the other system 2b will have reset and the communication problem with the target adaptor 4b will have been repaired or taken off-line. Thus, after an alert is received (at block 104) that the windows timeout timer expired, the monitoring ends and any processes or timers started in response to receiving an I/O request at block 104 are terminated (at block 105) because the communication problem in the system 2b including the target adaptor 4b is in the process of being repaired. During the system 2b error recovery procedure, the target adaptor 4b would go off-line and not affect loop operations.

[0021] At block 106, upon receiving an Input/Output (I/O) request, the device driver 22a starts (at block 108) an I/O delay timer for the received I/O request and forwards the I/O request to the adaptor 4a to execute against the target storage device 8a, 8b....8j. In certain implementations, the I/O delay timer reflects a maximum tolerable I/O delay dependent on the environment in which the storage system is deployed. Thus, the I/O delay timer in high availability systems is relatively shorter than systems requiring less availability. For instance, in high availability systems, the windows timeout period may run for several minutes to ensure the initiation of the error recovery process at the remote system 2b and the I/O delay period may be only a couple of seconds. Upon receiving (at block 110) notification that the monitored I/O request completed, the device driver 22a would stop (at block 112) the I/O delay timer that was started for the completed I/O request and control returns to block 106 to await any further received I/O requests or one of the other events at blocks 104 and 114. Upon receiving (at block 114) notification that one I/O delay timer expired, the device driver 22a issues (at block 116) a command

to the local adaptor 4a to get the identifier (ID) of the target adaptor 4b. As mentioned, in certain implementations, the local adaptor 4a could access the target adaptor 4b from a loop registry. If there is no ID for the target adaptor 4b in the loop registry, then that target adaptor 4b is not available on the loop. However, the unavailable target adaptor 4b may still affect I/O operations of the source adaptor 4a, such as by preventing the loop from completing a configuration. The target adaptor 4b may be only temporarily unavailable. If (at block 118) the get ID operation succeeded, i.e., the ID of the target adaptor 4b is included in the loop registry, then the device driver 22a issues (at block 120) a reset request to the received ID of the target adaptor 4b to cause a reset of the target adaptor 4b. Otherwise, if (at block 118) the get ID request did not succeed, then the device driver 22a returns to block 116 to issue additional get ID requests until the ID is received or the window timeout timer expires at block 104. The device driver 22a may be unable to obtain the ID if the system 2b including the target adaptor 4b, or the I/O bays 52a and 52b in the case of FIG. 2, is in the process of being reset.

15 [0022] In alternative implementations, the source adaptor 4a may query the target adaptor 4b over the loop 6 for the adaptor 4b ID. With the logic of FIG. 3, if the target adaptor 4b becomes available after one or more get ID requests fail at block 118 and before the windows timeout timer expires at block 104, then the device driver 22a would send the reset command to the adaptor 4b to cause the adaptor 4b to reset. If, during the reset operation, the target adaptor 4b is unable to establish communication with the system 2b, then the target adaptor 4b would not come back online and would not be available on the loop 6 until after the system 2b resets or power cycles and corrects the communication problem with the target adaptor 4b. If the target adaptor 4b becomes unavailable as a result of the reset, then the source adaptor 4a and any other adaptors on the loop 6 can continue operating on the loop 6 without concern that the problems at the off-line target adaptor 4b will cause problems.

25 [0023] In the described implementations, the device drivers 22a and 22b communicate an adaptor reset command. In alternative implementations, a system reset or power cycle command may be communicated. In implementations such as FIG. 2 where the two

adaptors are housed in the same system, then a problem may occur at the I/O bay 52a and 52b level including the adaptor, such that the adaptor, e.g., adaptor 54b, cannot communicate with the system even though the adaptor 54b can communicate on the loop. In such case, the device driver 72a detecting the problem in the I/O bay in the same system 50 may perform the reset with respect to the specific I/O bay 52b where the error is detected to avoid resetting or rebooting the entire system 50.

[0024] With the logic of FIG. 3, the device drivers 22a, 22b upon detecting errors at another peer adaptor will begin monitoring I/O requests to determine if any unacceptable I/O delays have occurred, i.e., I/O delays that extend the duration of the I/O delay timer. If no such unacceptable I/O delays occur, then the device drivers 22a, 22b would continue performing I/Os and the system 2a, 2b housing the target adaptor would eventually reset without affecting I/O processing at the other adaptor 4a, 4b. However, if an I/O request is delayed for the deemed unacceptable period, i.e., the I/O delay period, then a reset request will be sent to the adaptor 4b including the target adaptor 4b under the assumption that the problems in the system 2b including the target adaptor 4b are affecting the target adaptor's 4b ability to communicate on the loop 6. In cases where the device driver 22a, 22b cannot communicate with the target adaptor, i.e., cannot obtain the ID of the target at blocks 116-120, then the device driver 22a will continue to attempt to obtain the target adaptor 4b ID by resending the get ID request until either the identifier is received so that a reset request can be transmitted at block 120 or the window timeout timer expires at block 104 indicating that the system 4b housing the detected adaptor 4b with the communication problem has had sufficient time to initiate and complete an error recovery procedure.

[0025] The logic of FIG. 3 thus provides a greater degree of intelligent monitoring in cases where there is a problem in the system including one adaptor that will eventually result in a reset after a period of time that is longer than an acceptable I/O delay time given the availability requirements of the system. With the described implementations, unacceptable I/O delays are not experienced while waiting for the system or I/O bay including the other adaptor to reset. The described implementations are particularly

suited for situations where the detected problem between the target adaptor 4b and system 2b including the adaptor is not likely to prevent I/O requests at other adaptors from completing. The logic of FIG. 3 would allow one adaptor to continue processing I/O requests and wait for the detected adaptor to internally reset without having to interrupt I/O processing.

[0026] FIGs. 4a and 4b illustrate logic implemented in the device drivers 22a and 22b to account for the fact that the adaptor 4a, 4b in the system where the communication error is detected includes the master adaptor for the loop 6. The logic of FIGs. 4a and 4b provides for a modification of the logic of FIG. 3 to reduce any delays associated with switching the master if the master adaptor is in the target adaptor that will be subject to reset. The steps at block 200, 202, 204, 205, 206, 208, 210, 212, and 214 are identical to the steps performed at blocks 100, 102, 104, 105, 106, 108, 110, 112, and 114 in FIG. 3. The logic diverges in FIGs. 4a and 4b to provide for switching the master if necessary to avoid any delays in altering the master configuration. Upon receiving notification that an I/O timer expired (at block 214), the device driver 22a performs the steps at blocks 216, 218, and 220 (which are identical to steps 116, 118, and 120 in FIG. 3) to issue the command to do the reset of the system/I/O bay where the error was detected. As discussed, the reset may comprise a reset of the adaptor 4b only, or in additional implementations a power cycle command for the system 2 or power boundary including the target adaptor 4b.

[0027] Upon detecting the error at block 200, and concurrently with starting the window timeout timer (at block 202) and monitoring for received I/O requests (at block 206), an additional process is initiated at block 230 (FIGs. 4a and 4b). With respect to FIG. 4b, if (at block 230) the adaptor 4b (housed in the system 2b where the error was detected) is the master, then a timer is started (at block 232) for a master switch timeout period. The master switch timeout period is a timeout period selected to change the master in a manner that avoids any master switching delays should the target adaptor 4b in the system where the error was detected be reset. In certain implementations, the master switch timeout period is set to a value less than the window timeout timer and the

I/O delay timeout to ensure that the master is switched before the reset occurs following the I/O timer or window timer expiring, so that the resetting process will not be further delayed by having to wait for a switch of the master to another adaptor when the target adaptor 4b goes offline during the reset period. In this way, any I/O processing delays
5 due to a master switch operation are avoided.

[0028] In implementations where the I/O delay timeout is greater than the master switch timeout period, by the time the I/O delay timer expires, the master switch timer would have expired and triggered a change of master operation at block 230, 232, 234, and 236 if the target adaptor 4b subject to the reset is the master. Otherwise, if the I/O
10 delay timer is not greater than the master switch timer, then the reset process initiated in response to the I/O delay timer expiring at block 214 would further reset the master as part of the reset process.

[0029] Upon receiving (at block 234) an alert that the master switch timer expired, the device driver 22a issues (at block 236) a command to change the master to an adaptor
15 having a next highest unique identifier (ID) port, e.g., adaptor 4a, other than the adaptor 4b subject to the reset. After block 236, the process started at block 230 to switch the master ends. However, the process to monitor for any I/O requests and the process waiting for the window timeout timer would remain pending to perform any further steps according to the logic described in FIG. 4a.

[0030] The logic of FIGs. 4a and 4b provides additional checking over the logic of FIG. 3 to avoid any delays incurred when changing the mastership of the loop 6 from the target adaptor 4b being reset to another adaptor on the loop 6. In such case, the described implementations provide a master switch timeout period to ensure that the mastership of the loop 6 is switched before the target adaptor 4b is subject to a reset command to avoid
25 any I/O delays related to switching the mastership on the loop 6.

Additional Implementation Details

[0031] The preferred embodiments may be implemented as a method, apparatus or article of manufacture using standard programming and/or engineering techniques to

produce software, firmware, hardware, or any combination thereof. The term “article of manufacture” as used herein refers to code or logic implemented in hardware logic (e.g., an integrated circuit chip, Field Programmable Gate Array (FPGA), Application Specific Integrated Circuit (ASIC), etc.) or a computer readable medium (e.g., magnetic storage medium (e.g., hard disk drives, floppy disks,, tape, etc.), optical storage (CD-ROMs, optical disks, etc.), volatile and non-volatile memory devices (e.g., EEPROMs, ROMs, PROMs, RAMs, DRAMs, SRAMs, firmware, programmable logic, etc.). Code in the computer readable medium is accessed and executed by a processor. The code in which preferred embodiments are implemented may further be accessible through a transmission media or from a file server over a network. In such cases, the article of manufacture in which the code is implemented may comprise a transmission media, such as a network transmission line, wireless transmission media, signals propagating through space, radio waves, infrared signals, etc. Of course, those skilled in the art will recognize that many modifications may be made to this configuration without departing from the scope of the present invention, and that the article of manufacture may comprise any information bearing medium known in the art.

[0032] In the described implementations, the monitoring logic was implemented in the adaptor device driver, which is an operating system program. Alternatively, the monitoring logic of FIG. 3 may be implemented within the adaptor hardware.

[0033] In the described implementations, the detected error condition which triggered the logic of FIG. 3 indicated a problem in the system or I/O bay including the adaptor, where the adaptor itself is functioning properly. However, the logic of FIG. 3 may be initiated in response to other errors, such as internal adaptor hardware errors, line fault errors, ACK time out error in not receiving the required acknowledgment from the target adaptor, loss of synchronization error, code violation error in decoding a character that is not data or is a special character, protocol error, sequence error, frame reject error, etc. Further details of the errors that may trigger the error recovery logic of FIG. 3 are described in the publication "Serial Storage Architecture - Transport Layer 2 (SSA-

TL2)", Revision 5b, Working Draft T10.1/1147D (ANSI, April 4, 1997), which publication is incorporated herein by reference in its entirety.

[0034] In the described implementations, the reset sent by the device drivers 22a and 22b is a reset to the adaptor and not other components in the system housing the adaptor.

- 5 In alternative implementations, the reset may cause a reset or power cycle of the entire system housing the adaptor. In implementations where the system including the adaptor where the error is detected includes multiple power boundaries, i.e., multiple I/O bays, then the reset may cause a power cycle to the specific power boundary, i.e., I/O bay, including the target adaptor 4b where the problem is detected. Such a power boundary
- 10 specific reset would not affect other system components, such as other adaptors, on other power boundaries, such as other I/O bays.

- [0035]** In the described implementations, the storage devices and adaptors are all connected on a loop topology, using a loop protocol such as SSA, the Fibre Channel arbitrated loop protocol, etc. In alternative implementations, the two adaptors may be
- 15 connected to the storage devices in a non-loop topology. In the topology shown in FIG. 1, each adaptor may directly communicate with any of the storage devices 8a, 8b...8j. In additional implementations, one adaptor may have to communicate through another adaptor to reach a target storage device.

- [0036]** In certain implementations, the disk drives 8a, 8b...8j comprised magnetic hard
- 20 disk drives. In alternative implementations, the storage devices 8a, 8b...8j may comprise any storage device known in the art, such as optical disks, tapes, etc.

- [0037]** The preferred logic of FIGs. 3, 4a, and 4b describes specific operations occurring in a particular order. In alternative implementations, certain of the logic operations may be performed in a different order, modified or removed. Moreover, steps
- 25 may be added to the above described logic and still conform to the described implementations. Further, operations described herein may occur sequentially or certain operations may be processed in parallel.

[0038] The foregoing description of the preferred embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be

5

10

1927-1928 1929-1930 1931-1932 1933-1934 1935-1936 1937-1938 1939-1940 1941-1942 1943-1944 1945-1946 1947-1948 1949-1950 1951-1952 1953-1954 1955-1956 1957-1958 1959-1960 1961-1962 1963-1964 1965-1966 1967-1968 1969-1970 1971-1972 1973-1974 1975-1976 1977-1978 1979-1980 1981-1982 1983-1984 1985-1986 1987-1988 1989-1990 1991-1992 1993-1994 1995-1996 1997-1998 1999-2000 2001-2002 2003-2004 2005-2006 2007-2008 2009-2010 2011-2012 2013-2014 2015-2016 2017-2018 2019-2020 2021-2022 2023-2024 2025-2026 2027-2028 2029-2030 2031-2032 2033-2034 2035-2036 2037-2038 2039-2040 2041-2042 2043-2044 2045-2046 2047-2048 2049-2050 2051-2052 2053-2054 2055-2056 2057-2058 2059-2060 2061-2062 2063-2064 2065-2066 2067-2068 2069-2070 2071-2072 2073-2074 2075-2076 2077-2078 2079-2080 2081-2082 2083-2084 2085-2086 2087-2088 2089-2090 2091-2092 2093-2094 2095-2096 2097-2098 2099-2100 2101-2102 2103-2104 2105-2106 2107-2108 2109-2110 2111-2112 2113-2114 2115-2116 2117-2118 2119-2120 2121-2122 2123-2124 2125-2126 2127-2128 2129-2130 2131-2132 2133-2134 2135-2136 2137-2138 2139-2140 2141-2142 2143-2144 2145-2146 2147-2148 2149-2150 2151-2152 2153-2154 2155-2156 2157-2158 2159-2160 2161-2162 2163-2164 2165-2166 2167-2168 2169-2170 2171-2172 2173-2174 2175-2176 2177-2178 2179-2180 2181-2182 2183-2184 2185-2186 2187-2188 2189-2190 2191-2192 2193-2194 2195-2196 2197-2198 2199-2200 2201-2202 2203-2204 2205-2206 2207-2208 2209-2210 2211-2212 2213-2214 2215-2216 2217-2218 2219-2220 2221-2222 2223-2224 2225-2226 2227-2228 2229-2230 2231-2232 2233-2234 2235-2236 2237-2238 2239-2240 2241-2242 2243-2244 2245-2246 2247-2248 2249-2250 2251-2252 2253-2254 2255-2256 2257-2258 2259-2260 2261-2262 2263-2264 2265-2266 2267-2268 2269-2270 2271-2272 2273-2274 2275-2276 2277-2278 2279-2280 2281-2282 2283-2284 2285-2286 2287-2288 2289-2290 2291-2292 2293-2294 2295-2296 2297-2298 2299-2300 2301-2302 2303-2304 2305-2306 2307-2308 2309-2310 2311-2312 2313-2314 2315-2316 2317-2318 2319-2320 2321-2322 2323-2324 2325-2326 2327-2328 2329-2330 2331-2332 2333-2334 2335-2336 2337-2338 2339-2340 2341-2342 2343-2344 2345-2346 2347-2348 2349-2350 2351-2352 2353-2354 2355-2356 2357-2358 2359-2360 2361-2362 2363-2364 2365-2366 2367-2368 2369-2370 2371-2372 2373-2374 2375-2376 2377-2378 2379-2380 2381-2382 2383-2384 2385-2386 2387-2388 2389-2390 2391-2392 2393-2394 2395-2396 2397-2398 2399-2400 2401-2402 2403-2404 2405-2406 2407-2408 2409-2410 2411-2412 2413-2414 2415-2416 2417-2418 2419-2420 2421-2422 2423-2424 2425-2426 2427-2428 2429-2430 2431-2432 2433-2434 2435-2436 2437-2438 2439-2440 2441-2442 2443-2444 2445-2446 2447-2448 2449-2450 2451-2452 2453-2454 2455-2456 2457-2458 2459-2460 2461-2462 2463-2464 2465-2466 2467-2468 2469-2470 2471-2472 2473-2474 2475-2476 2477-2478 2479-2480 2481-2482 2483-2484 2485-2486 2487-2488 2489-2490 2491-2492 2493-2494 2495-2496 2497-2498 2499-2500 2501-2502 2503-2504 2505-2506 2507-2508 2509-2510 2511-2512 2513-2514 2515-2516 2517-2518 2519-2520 2521-2522 2523-2524 2525-2526 2527-2528 2529-2530 2531-2532 2533-2534 2535-2536 2537-2538 2539-2540 2541-2542 2543-2544 2545-2546 2547-2548 2549-2550 2551-2552 2553-2554 2555-2556 2557-2558 2559-2560 2561-2562 2563-2564 2565-2566 2567-2568 2569-2570 2571-2572 2573-2574 2575-2576 2577-2578 2579-2580 2581-2582 2583-2584 2585-2586 2587-2588 2589-2590 2591-2592 2593-2594 2595-2596 2597-2598 2599-2600 2601-2602 2603-2604 2605-2606 2607-2608 2609-2610 2611-2612 2613-2614 2615-2616 2617-2618 2619-2620 2621-2622 2623-2624 2625-2626 2627-2628 2629-2630 2631-2632 2633-2634 2635-2636 2637-2638 2639-2640 2641-2642 2643-2644 2645-2646 2647-2648 2649-2650 2651-2652 2653-2654 2655-2656 2657-2658 2659-2660 2661-2662 2663-2664 2665-2666 2667-2668 2669-2670 2671-2672 2673-2674 2675-2676 2677-2678 2679-2680 2681-2682 2683-2684 2685-2686 2687-2688 2689-2690 2691-2692 2693-2694 2695-2696 2697-2698 2699-2700 2701-2702 2703-2704 2705-2706 2707-2708 2709-2710 2711-2712 2713-2714 2715-2716 2717-2718 2719-2720 2721-2722 2723-2724 2725-2726 2727-2728 2729-2730 2731-2732 2733-2734 2735-2736 2737-2738 2739-2740 2741-2742 2743-2744 2